REMARKS

The Examiner's action dated December 14, 2004, has been received, and its contents carefully noted.

The Examiner has rejected original claims 1 to 14 under 35 U.S.C. §103(a) as being unpatentable over Cheriton *et al.* (US Patent No. 6,675,200) in view of Grun (US Patent No. 6,272,591).

This rejection is respectfully traversed on two distinct grounds.

First, the present application claims priority from a US provisional application 60/215,994 dated July 5, 2000. Cheriton *et al.* has a filing date of May 10, 2000, i.e. less than two months before applicants' filing date. Applicants respectfully notes that the invention *as claimed at least in original claim 6, amended claims 7, 8, 9, 10, 11, 12, and new claims 15 to 21* was conceived before May 10, 2000 and applicants and their attorneys proceeded diligently to prepare and file the provisional application. In consequence Cheriton *et al.* is not available as prior art against the claims based on the disclosure in the provisional application. A Declaration under 37 CFR 1.131 signed by the inventors is enclosed together with evidence in the form of documents identified in the Declaration that show conclusively that the invention was conceived prior to Cheriton *et al.* i.e. before May 10, 2000 and diligence was exercised to file the provisional application, given the short period that elapsed between May 10 and July 5, 2000.

Secondly, Applicant disputes the Examiner's analysis in Section 4 of the Office Action where, we believe, he mistakenly avers that Cheriton teaches a reading device having an "RDMA engine

each adapted to receive data packets associated with a transaction on a respective communication channel." Indeed, this statement appears to be irreconcilable with the Examiner's subsequent statement in the next paragraph, with which we concur, that Cheriton *et al.* does not discuss a plurality of RDMA engines. Since Cheriton *et al.* does not discuss a plurality of RDMA engines there appears to be basis to associating a "*respective* communication channel" with the RDMA engine. Claim 7 clearly requires that there be multiple communication channels each associated with a respective RDMA engine and there appears nothing in Cheriton *et al.* to suggest that data is conveyed on more than a single communication channel. Moreover, nor does Cheriton *et al.* deal at all with the problem of dividing a large data transfer to be sent in pieces simultaneously over multiple RDMA engines. Thus, claim 1, which is not limited to data being conveyed along more than a single communication channel, is also distinguished over Cheriton *et al.*

Cheriton *et al.* appears to be directed to the problem of writing data packets directly to memory when using the TCP/IP (Internet) protocol without the need to make copies of large blocks of data. Applicant has carefully reviewed the sections of Cheriton *et al.* on the basis of which the Examiner reaches his conclusions, but can find nothing to suggest that Cheriton *et al.* even considers the problem to whose solution the claims of the present application are directed, namely conveying data to multiple memory units so as to allow a receiving device to determine when a complete data transfer has arrived, while avoiding the need for interrupts to be generated by the RDMA for each data packet.

That the present invention is directed to reducing interrupts is clear from the description on page 4, lines 7-13 of the application and also at the end of the first page of the original provisional application, both of which read *essentially* as follows:

> "This solution has the undesirable condition that it results in an interrupt being generated for each packet. The receiver is interested in knowing when the entire transaction (comprising all the data packets) has completed, and all the extra interrupts/ callbacks for the small data transfers consume resources that could otherwise be used for other purposes."

It is to be noted that this feature finds explicit reference in the amended claims. For example, clause (b) of claim 1 now recites:

> (b) for each memory unit receiving said respective final data packet generating an interrupt **such that a single interrupt is generated for each memory unit receiving data** and informing the receiving device how much data was received by said memory unit,

Corresponding amendments have been effected to the remaining independent claims.

Grun does use multiple channels to accomplish a large data transfer. However, in Grun, as shown in Fig. 2 and described in detail in col. 3, the RAID device itself sets up the different channels, specifies the amount of (and which) data to be sent over each channel, and checks when each of the mini-transfers is complete. The relevant description of Grun reads:

> "FIG. 2 is a flowchart of the steps executed by RAID device 40 in one embodiment of the present invention when an I/O request is received from host computer 10 to store a data block in RAID device 40. It is assumed that RAID device 40 stripes data blocks across "N" disk drives.

> "The request is received at step 110 and includes the location in memory 16 of host computer 10 where the data block is stored. Driver 14 stores the I/O request in a location of memory 16. In accordance with the VI specification, driver 14 posts a descriptor that refers to the I/O request (i.e., specifies the location in memory 16 where the I/O request is stored) to a send queue in transport 20. Driver 14 then rings a doorbell in NIC 25. The doorbell tells NIC 25 to look in the send queue for the descriptor. NIC 25 then fetches the descriptor and performs the task. The task places an I/O request message on connection 30 to be transmitted. The receiving device (i.e., RAID device 40) of the I/O request also has a NIC (i.e., NIC 42) that receives the I/O request message from connection 30.
>
> "The I/O request message contains information specifying the location in host memory 16 from which the data is to be moved, and specifies where in RAID device 40 the data is to be stored. The location in host memory 16 is specified with a virtual address memory handle pair in accordance with the VI specification. RAID device 40 uses the information contained in the I/O request message to build descriptors to accomplish the actual data movement from host computer 10 to RAID device 40. For example, in response to receiving the request from host computer 10, in one embodiment RAID device 40 initiates a data transfer from host computer 10. The data transfer is initiated using a VI Remote Direct Memory Access ("RDMA") transfer facility. " [col. 3, lines 14-46]

It thus emerges that in Grun, the RAID device initiates data transfer from specified locations in host memory and so *knows* when the data is completely transferred. Hence the very problem to whose solution the present invention is directed is avoided in Grun.

Specifically, in the present invention, a first entity i.e. the sending machine, decides on the sizes of the mini-transfers to be sent over the various channels, and a different entity i.e. the receiving machine, must determine when the complete data

transfer has arrived based on information in the mini-transfers. This solves the problem to which the present invention is directed, namely enabling the receiving machine to know from the packet information alone when the job is complete across all the mini-transfers, even though the receiving machine did not specify which (and how much) data is to be transferred over each communication channel. This idea is not at all addressed in Grun, and is not derivable even from the combination of the Cheriton and Grun patents.

Moreover, since Grun does not even address the problem with which the present invention is concerned there is no reason for one skilled in the art to combine Grun with Cheriton *et al*.

The assertion of the "motivation" for combining the reference teachings, presented in the explanation of the rejection of claim 7, does not represent a valid reason for combining the references because no evidence has been presented that it would be known in the art that the proposed combination would provide high bandwidth message-passing. In the absence of such evidence, it can only be concluded that this advantage is evidence of non-obviousness.

Favorable reconsideration and allowance are respectfully requested.

If the above amendment should not now place the application in condition for allowance, the Examiner is invited to call undersigned counsel to resolve any remaining issues.

Respectfully submitted,

BROWDY AND NEIMARK, P.L.L.C.
Attorneys for Applicant

By _____
Jay M. Finkelstein
Registration No. 21,082

JMF:dtb
Telephone No.: (202) 628-5197
Facsimile No.: (202) 737-3528
G:\BN\C\cohn\Satran1\pto\AMD 19APR05.doc